
Unifying Subsampling Pattern Variations for Compressed Sensing MRI with Neural Operators

Armeet Singh Jatyani* Jiayun Wang* Zihui Wu
Miguel Liu-Schiaffini Bahareh Tolooshams Anima Anandkumar
Department of Computing and Mathematical Sciences, Caltech
{armeet, peterw, zwu2, mliuschi, btoloosh, anima}@caltech.edu

Abstract

Compressed Sensing MRI (CS-MRI) reconstructs images of the body’s internal anatomy from undersampled and compressed measurements, thereby reducing scan times and minimizing the duration patients need to remain still. Recently, deep neural networks have shown great potential for reconstructing high-quality images from highly undersampled measurements. However, since deep neural networks operate on a fixed discretization, one needs to train multiple models for different measurement subsampling patterns and image resolutions. This approach is highly impractical in clinical settings, where subsampling patterns and image resolutions are frequently varied to accommodate different imaging and diagnostic requirements. We propose a unified model that is robust to different subsampling patterns and image resolutions in CS-MRI. Our model is based on neural operators, a discretization-agnostic architecture. We use neural operators in both image and measurement (frequency) space, which capture local and global image features for MRI reconstruction. Empirically, we achieve consistent performance across different subsampling rates and patterns, with up to 4x lower NMSE and 5 dB PSNR improvements over the state-of-the-art method. We also show the model is agnostic to image resolutions with zero-shot super-resolution results. Our unified model is a promising tool that is agnostic to measurement subsampling and imaging resolutions in MRI, offering significant utility in clinical settings where flexibility and adaptability are essential for efficient and reliable imaging.

1 Introduction

Magnetic Resonance Imaging (MRI) [3, 4] is a method of obtaining high-quality images of the body’s internal anatomy by using magnetic fields and radio waves. MRI is a popular non-invasive imaging technology, used in numerous medical and scientific applications such as neurosurgery [5], clinical oncology [6], diagnostic test [7], neuroscience [8], and pharmaceutical research [9]. Despite its popularity, MRI is greatly limited by a slow data acquisition process, which sometimes requires patients to remain still for an hour [10, 11]. Hence, accelerating MRI scan times has garnered tremendous attention [12, 13, 14, 15, 1].

Compressed Sensing (CS) [16], which revolutionized the recovery of signals from compressed data, utilizes sparse coding [17] to enable MR imaging at sub-Nyquist subsampling rates, significantly reducing the number of required measurements [18]. In practice, Compressed Sensing MRI (CS-MRI) reduces acquisition time, substantially increasing the clinical utility of the imaging method. In this context, MRI can be formulated as an ill-posed inverse problem [19, 20, 1], where prior information about MR images is essential for successful reconstruction. CS-MRI can also be viewed as image

*Equal contribution.

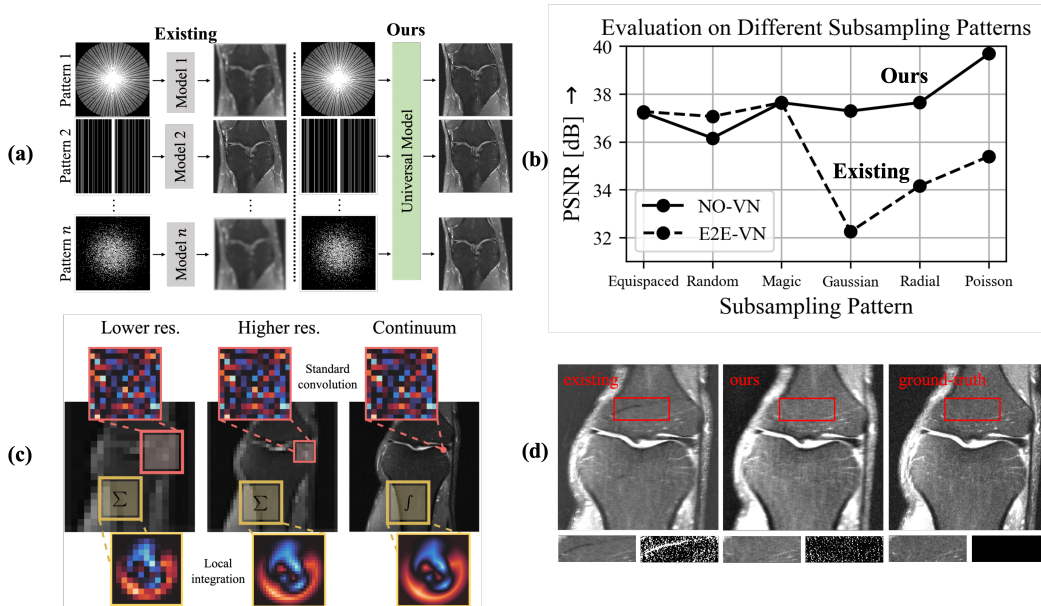


Figure 1: (a) We propose a unified model for different CS-MRI reconstruction settings. Existing models [1, 2] operate on a specific subsampling pattern, and require training additional models for additional patterns. We propose a single discretization-agnostic model that operates on multiple subsampling patterns for CS-MRI. (b) Our model (NO-VN) outperforms SOTA in cross-pattern generalization. Notably, our model achieves consistent performance across different subsampling patterns. (c) Comparison between our DISCO local integration operator (yellow) and the standard 2d convolution (red). As resolution increases to continuum, our DISCO local operator converges to a local integral, whereas the standard convolutional kernel converges to a point-wise operator in image space. (d) Our model outperforms SOTA (E2E-VN) on zero-shot super-resolution. Last row: zoomed-in visualization and difference from the ground truth.

recovery from compressed data [13, 21]. CS [22, 13, 23] achieves CS-MRI by assuming a sparsity prior on the image in a transform domain (e.g., discrete wavelet transform [24]). While sparse priors have been quite successful in compressing data [25, 26], recent work goes beyond classical priors and proposes learning the underlying data structures using deep learning [14, 1, 27]. Current state-of-the-arts [1, 15] follow this approach to find an *end-to-end* mapping from the undersampled measurements to the image reconstruction while learning implicit priors directly in the image space [1, 27] or the frequency domain [28, 29]. While successful, such methods only work for a fixed discretization (resolution) of the measurement data and output image. On the other hand, in clinical settings, radiologists need to constantly change resolutions for imaging and diagnosis purposes, which requires many machine learning models. A unified model that is agnostic to discretizations would greatly improve efficiency.

Neural operators are a class of deep learning architectures designed to learn maps between infinite-dimensional function spaces. As such, neural operators are discretization and resolution agnostic, and thus well equipped for accelerated MRI, where subsampling settings can vary in pattern, rate, and resolution. In this work, we make the following contributions.

1. We propose a unified model that is robust to different subsampling patterns and image resolutions in CS-MRI (Fig. 1a). We learn reconstruction priors in function space, which is discretization-agnostic via neural operators (Fig. 1c). Our model follows an unrolled variational network design [15, 1], incorporating neural operators (NOs) that operate on both the measurement and image space. The measurement space neural operator makes our framework agnostic to different measurement subsampling patterns, and the image space neural operator makes the framework agnostic to different image resolutions. This approach allows the model to capture both local and global features of images, leveraging the duality of the image and frequency spaces. To the best of our knowledge, this is the first application of neural operators to MRI.
2. We empirically demonstrate that our model is robust to different measurement subsampling rates and patterns (Fig. 1a). Our model performs consistently across these variations, whereas the

existing existing method drops in performance (Fig. 1b). We achieve up to 4x lower NMSE and 5 dB PSNR improvement from the baseline when evaluating on different subsampling patterns.

3. Finally, we show that our model outperforms the state-of-the-art in zero-shot super-resolution inference (Fig. 1d)

2 Related Works

Accelerated MRI. One way to accelerate MRI scan speed is parallel imaging, in which multiple receiver coils acquire different views of the objective interest simultaneously, and then combine them into a single image [12, 30, 31]. Paired with compressed sensing (CS-MRI), pre-defined priors or regularization filters can be leveraged to improve reconstruction quality [22, 13]. Recent works have shown that learned deep-learning priors outperform hand-crafted priors in reconstruction fidelity. Convolutional neural networks (CNNs) [14, 15, 32, 1], variational networks [15, 1], and generative adversarial networks (GANs) [14, 33] have all demonstrated superior performance in CS-MRI reconstruction from subsampled measurements. However, unlike conventional compressed sensing which operates in the function space, the aforementioned deep learning methods operate on a fixed resolution. As a result, changes in resolution lead to degradation in performance.

Discretization-Agnostic Learning and Neural Operators. Empirically, researchers have shown that diffusion models have relatively consistent performance with different measurement subsampling patterns in accelerated MRI [34]. However, diffusion models usually take more runtime at inference. Additionally, they are not fundamentally discretization-agnostic by design. Neural operators, as introduced by [35, 36], are deep learning architectures specifically designed to learn mappings between infinite-dimensional function spaces. They are discretization-agnostic, allowing evaluation at any resolution, and converge to a desired operator as the resolution approaches infinity. Neural operators have empirically achieved good performance across various applications, such as material science [37], weather forecasting [38], and photoacoustic imaging [39]. Neural operators have various designs. For example, the Fourier neural operator (FNO) [40], which performs global convolutions, has shown consistent discretization-agnostic performance in various applications [35]. Other designs of neural operators [41, 42] rely on integration with locally-supported kernels to capture local features, which has shown to be useful in applications where local features are important, such as modeling turbulent fluids. Additionally, neural operators with local integrals can be made efficient with parallel computing compared to those requiring global integrals. We use neural operators to build our MRI framework that is agnostic to subsampling patterns and image resolutions.

3 Methods

3.1 CS-MRI with Unrolling

Background. In MRI, images \mathbf{x} depicting patient anatomy are reconstructed by acquiring measurements \mathbf{k} , in the frequency domain. They are related as

$$\mathbf{k} = \mathcal{F}(\mathbf{x}) + \epsilon \tag{1}$$

where ϵ is the measurement noise and \mathcal{F} is the Fourier transformation. In this paper, we consider multi-coil MRI [43, 44] data, which images different regions of the anatomy in parallel. The forward process of the i^{th} coil measures $\mathbf{k}_i = \mathcal{F}(S_i \mathbf{x}) + \epsilon_i$ where S_i is a position-dependent sensitivity map for the i^{th} coil. To speed up the imaging process, measurements are subsampled as $\tilde{\mathbf{k}} = M\mathbf{k}$, where M is a binary mask that selects a subset of the \mathbf{k} -space points. Classical compressed sensing methods reconstruct the image $\hat{\mathbf{x}}$ by solving an optimization problem

$$\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x}} \frac{1}{2} \sum_i \left\| \mathcal{A}(\mathbf{x}) - \tilde{\mathbf{k}} \right\|^2 + \lambda \Psi(\mathbf{x}) \tag{2}$$

where i is coil index, $\mathcal{A}(\cdot) = MFS(\cdot)$ is the linear forward operator, and $\Psi(\mathbf{x})$ is a prior regularization term. The optimization objective can be considered as a combination of physics constraint and prior. While the above optimization can be solved using classical optimization toolboxes, an increasing line of works uses deep learning for MRI [45, 1]. Among them, unrolled networks [15, 1] have gained popularity as they incorporate the known forward model, resulting in state-of-the-art performance. Unrolling, which started with the nominal work of LISTA [46], proposes to design




Figure 2: **MRI reconstruction pipeline.** We learn priors in function space with neural operators in the measurement space KNO and image space INO.

networks using iterations of an optimization algorithm to solve inverse problems. This approach incorporates domain knowledge (i.e., the forward model) and leverages deep learning to learn implicit priors from data [47, 48]. In the context of MRI and assuming a differential regularization term, the optimization problem is expanded to iterative gradient descent steps with injected CNN-based data priors. Each layer mimics the gradient descent step from \mathbf{x}^t to \mathbf{x}^{t+1} :

$$\mathbf{x}^{t+1} = \mathbf{x}^t - \eta^t A^*(A(\mathbf{x}^t) - \tilde{\mathbf{k}}) + \lambda^t \text{CNN}(\mathbf{x}^t) \quad (3)$$

where η^t controls the weight of *data consistency* term and λ^t controls that of the data-driven prior term. The data consistency term samples the data in the frequency domain, hence it is applicable to any spatial resolution. However, the prior term only operates on a specific resolution with CNNs. This means when changing the subsampling patterns, one needs another CNN trained for that setting, which greatly limits the flexibility of the reconstruction system.

Extending to Neural Operators. We learn the prior in function space via discretization-agnostic neural operators KNO and INO. First, we use a measurement space neural operator KNO to learn $\mathbf{k}^0 = \text{KNO}(\tilde{\mathbf{k}})$. Then, $\mathbf{x}^0 = \mathcal{F}^{-1}(\mathbf{k}^0)$ becomes the input to the unrolled layers, each of which features a data consistency loss and the image space INO for image prior learning:

$$\mathbf{x}^{t+1} = \mathbf{x}^t - \eta^t A^*(A(\mathbf{x}^t) - \tilde{\mathbf{k}}) + \lambda^t \text{INO}(\mathbf{x}^t) \quad (4)$$

Note that we follow existing works [45, 1] and do not include priors on \mathbf{k} in each unrolled layer anymore. Our framework flexibly works for different resolutions.

Framework Overview. Fig. 2 depicts the pipeline of our neural operator framework for MRI reconstruction. The subsampled measurement $\tilde{\mathbf{k}}$ is first fed to a neural operator KNO which operates in measurement \mathbf{k} space to learn global image features and then inverse Fourier transformed to get an image. Following Eqn. 4, we iterate a few cascades of unrolled layers, consisting of a neural operator INO which operates in image \mathbf{x} space and a data consistency update.

3.2 Neural Operator Designs

Neural operators learn mappings between function spaces and can be used for discretization-agnostic MRI reconstruction. Since accurate MRI reconstruction is a function of both local and global features of the image, we propose to incorporate both global and local inductive biases into our neural operator architecture. We first discuss how we learn local features with local integration operators.

Local Features with Local Integration Operator. Historically, the most common method of embedding a local inductive bias into deep neural networks has been by using locally-supported convolutional kernels, as in convolutional neural networks (CNNs). However, standard discrete convolutional kernels used in CNNs do not satisfy the resolution-agnostic properties of neural operators. Specifically, [42] shows that CNN-style convolutional kernels converge to pointwise linear operators as the resolution is increased, instead of the desired local integration in the limit of infinite resolution. For a kernel κ and input function v defined over some compact subset $D \subset \mathbb{R}^d$, the *local convolution operator* is given by

$$(k \star v)(y) = \int_D \kappa(x - y) \cdot v(x) dx. \quad (5)$$

Given a particular set of input points $(x_j)_{j=1}^m \subset D$ with corresponding quadrature weights q_j and output positions $y_i \in D$, [42] adopts the framework of discrete-continuous convolutions (DISCO) [49]

and approximates the convolution as

$$(k \star v)(y_i) \approx \sum_{j=1}^m \kappa(x_j - y_i) \cdot v(x_j) q_j. \quad (6)$$

[42] proposes to parameterize κ as a linear combination of pre-defined basis functions κ^ℓ : $\kappa = \sum_{\ell=1}^L \theta^\ell \cdot \kappa^\ell$, where θ^ℓ are learnable parameters. The convolutional kernel is thus parameterized by a finite number of parameters, independently of the grid on which the kernel is evaluated. Since we are operating on an equidistant grid on a compact subset of \mathbb{R}^2 , we follow [42] and implement Eqn. 6 using standard convolutional kernels (thus enjoying the benefits of acceleration on GPUs using standard deep learning libraries) with two crucial modifications: (1) the kernel itself is defined as a linear combination of basis functions κ^ℓ , and (2) the size of the kernel scales with the input resolution so as to remain a fixed size w.r.t. the input domain. We adopt the same basis functions as [42] in our experiments, and we use the local integration operator as the resolution-agnostic building block for the measurement space and image space operators.

Global Features. A common neural operator architecture for learning global features is the Fourier neural operator (FNO) [40]. FNO takes the Fourier transform of the input, truncates the result beyond some fixed number of modes, and pointwise multiplies the result with a learned weight tensor, which is equivalent to a global convolution on the input by the convolution theorem. Interestingly, the forward process of MRI is a Fourier transformation, which means that local operations in measurement \mathbf{k} space are equivalent to global operators in image \mathbf{x} space and vice versa, due to their duality. Following FNO, we could apply a pointwise multiplication between the measurement \mathbf{k} and a learned weight tensor to capture global image features. However, FNO truncates high frequencies, which are crucial for MRI reconstruction. To address this, we directly apply the DISCO local integration operator on the measurement space to capture global image features without feature map truncation.

UDNO: the Building Block. Without loss of generality, we make both the image neural operator INO and measurement neural operator KNO be a local neural operator that captures local features in the corresponding domain. As discussed earlier, such a design learns both global and local image features. Each operator consists of multiple sub-layers, which we refer to as the U-Shaped DISCO Neural Operator, or UDNO. U-shaped networks are one of the most popular architectures in computer vision and have shown great performance in various applications from medical imaging to diffusion [31, 50, 51]. The UDNO follows the encoder/decoder architecture of the U-Net [31], with regular convolutions being replaced by neural operator DISCO layers. Additional details are provided in Fig.4 and Section A of the appendix.

Loss. The parameters of the proposed neural operator are estimated from the training data by minimizing the structural similarity loss between the reconstruction \mathbf{x} and the ground truth image \mathbf{x}^*

$$\mathcal{L}(\hat{\mathbf{x}}, \mathbf{x}^*) = -\text{SSIM}(\hat{\mathbf{x}}, \mathbf{x}^*), \quad (7)$$

where SSIM is the Structural Similarity index [52].

4 Experiments

For efficiency, and to focus on comparing the choice of network architecture for learning reconstruction priors, rather than unrolling dynamics, we conduct multiple-pattern and rate experiments with single-cascade models.

4.1 Dataset and Setup

The **fastMRI Dataset** is a large and open dataset of knee and brain MRIs. We use their multi-coil knee reconstruction subset with 34,742 slices for training and 7,135 slices for evaluation [44].

Subsampling Patterns and Rates. We use equispaced, random, magic, Gaussian, radial, and Poisson subsampling patterns and 2x, 4x, 6x, and 8x subsampling rates (see Fig. 5). Higher rates result in sparser k -space samples and shorter imaging-times at the cost of a more difficult inversion process. Section B in the appendix provides additional subsampling details.

Neural Operator Model. Our neural operator model follows Fig. 2. The KNO (measurement space neural operator) and INO (image space neural operator) are implemented as UDNOs with 2 input and

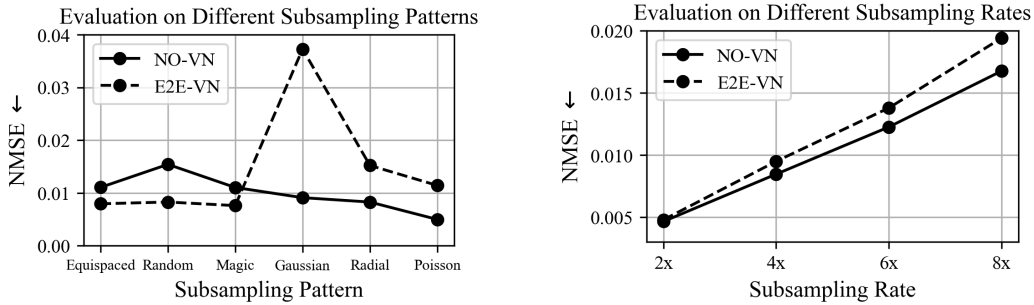


Figure 3: **Left:** Multi pattern results. We fix the subsampling rate to 2x, and evaluate both models across different patterns to measure cross-pattern generalization performance. **Right:** We fix the pattern to equispaced and vary the subsampling rate.

output channels. This is because complex numbers, commonly used in MRI data, are represented using two channels: one for the real part and one for the imaginary part. We train both our model and baseline with a total batch size of 16 across 4 A100 (40G) GPUs. We provide UDNO details, DISCO kernel basis configurations, and training hyper-parameters in Section B of the appendix.

SOTA Baseline. We compare with the SOTA End-to-End VarNet (E2E-VN) method [1], which shares a similar network structure with our approach, but uses traditional resolution-dependent CNNs.

Evaluation Protocols. We evaluate image reconstruction performance using normalized mean square error (NMSE), peak signal-to-noise ratio (PSNR), and structural similarity index measure (SSIM) which are standard for the fastMRI dataset and MRI [44].

4.2 Agnostic to Measurement Subsampling

We pre-train both our NO-VN and E2E-VN (SOTA) on 2x *equispaced* samples for 20 epochs. Result visualizations are provided in Section C of the appendix. Generally, we observe that the performance of our NO-VN model withstands variations in measurement subsampling pattern and acceleration rate, whereas SOTA performance degrades with variation.

Agnostic to Subsampling Patterns. Both pre-trained models are trained for an additional 20 epochs on a small dataset (3,474 samples) of Gaussian, radial, and Poisson 2x accelerated samples. Visualizations are provided in Section C of the appendix. For irregular patterns, we achieve an average improvement of 4.269 dB PSNR over the state-of-the-art baseline. On rectilinear patterns, our performance remains comparable to the baseline. Visualizations are in Fig. 7 of the appendix.

Agnostic to Subsampling Rate. Both pre-trained models are trained for an additional 20 epochs on a small dataset (3,474 samples) of 2x, 4x, 6x, 8x equispaced accelerated samples. Our model consistently outperforms the SOTA, achieving 0.409 dB higher PSNR on average. Visualizations are in Fig. 7 of the appendix.

Number of Cascades. In Section D of the appendix, we present cascade experiments to show that our architecture improves with more unrolling cascades.

4.3 Zero-Shot Super-Resolution

We pre-train our NO-VN and the E2E-VN (baseline) models on 320x320 samples. We then evaluate their performance on higher-resolution 640x640 samples. Since k-space interpolation leads to a larger field of view in image space, we only evaluate at super-resolution in image space (INO). Compared to our neural operator model, the convolution-based E2E-VN produces reconstructions with noticeable artifacts, which can be misleading in clinical settings (see Fig.1d). Ensuring the reliability and consistency of a model is crucial for accurate diagnosis and treatment.

5 Conclusion

We introduce a unified model for Compressed Sensing MRI (CS-MRI) that overcomes the need for training multiple deep learning models to handle different subsampling patterns and image resolutions, which is a common issue in clinical practice. By utilizing neural operators, which are agnostic to the input discretization, our model captures both local and global features for more flexible MRI reconstruction. On the fastMRI dataset, our unified model achieves consistent performance across varying subsampling patterns and outperforms state-of-the-art methods in terms of accuracy and robustness. The unified model also improves super-resolution tasks. In future work, we aim to explore super-resolution on higher-resolution MRI datasets, as we are currently limited by the resolution of the fastMRI dataset. Our approach offers a versatile solution for more efficient and reliable MRI in medical imaging, with significant utility in clinical settings where flexibility and adaptability to varying subsampling patterns and resolutions are crucial.

References

- [1] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated mri reconstruction. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23*, pages 64–73. Springer, 2020.
- [2] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical Image Analysis*, 80:102479, 2022.
- [3] G.A. Wright. Magnetic resonance imaging. *IEEE Signal Processing Magazine*, 14(1):56–66, 1997.
- [4] Vadim Kuperman. *Magnetic resonance imaging: physical principles and applications*. Elsevier, 2000.
- [5] V Seifert, M Zimmermann, C Trantakis, H-E Vitzthum, K Kühnel, A Raabe, F Bootz, J-P Schneider, F Schmidt, and J Dietrich. Open mri-guided neurosurgery. *Acta neurochirurgica*, 141:455–464, 1999.
- [6] Dow-Mu Koh and David J Collins. Diffusion-weighted mri in the body: applications and challenges in oncology. *American Journal of Roentgenology*, 188(6):1622–1635, 2007.
- [7] DJ Husband, KA Grant, and CS Romaniuk. Mri in the diagnosis and treatment of suspected malignant spinal cord compression. *The British journal of radiology*, 74:15–23, 2001.
- [8] Denis Le Bihan. Looking into the functional architecture of the brain with diffusion mri. *Nature reviews neuroscience*, 4(6):469–480, 2003.
- [9] J Craig Richardson, Richard W Bowtell, Karsten Mäder, and Colin D Melia. Pharmaceutical applications of magnetic resonance imaging (mri). *Advanced drug delivery reviews*, 57(8):1191–1209, 2005.
- [10] Yutong Chen, Carola-Bibiane Schönlieb, Pietro Liò, Tim Leiner, Pier Luigi Dragotti, Ge Wang, Daniel Rueckert, David Firmin, and Guang Yang. Ai-based reconstruction for fast mri—a systematic review and meta-analysis. *Proceedings of the IEEE*, 110(2):224–245, 2022.
- [11] Dilbag Singh, Anmol Monga, Hector L de Moura, Xiaoxia Zhang, Marcelo VW Zibetti, and Ravinder R Regatte. Emerging trends in fast mri using deep-learning reconstruction on undersampled k-space data: a systematic review. *Bioengineering*, 10(9):1012, 2023.
- [12] Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.
- [13] Michael Lustig, David L Donoho, Juan M Santos, and John M Pauly. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008.

- [14] Patricia M Johnson and Maria Drangova. Conditional generative adversarial network for 3d rigid-body motion correction in mri. *Magnetic resonance in medicine*, 82(3):901–910, 2019.
- [15] Hammernik K, Klatzer T, Kobler E, et al. Learning a variational network for reconstruction of accelerated MRI data, 2018.
- [16] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [17] David L Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via l_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- [18] Xi Chen, Zhuizhuan Yu, Sebastian Hoyos, Brian M Sadler, and Jose Silva-Martinez. A subnyquist rate sampling receiver exploiting compressive sensing. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 58(3):507–520, 2010.
- [19] Charles W Groetsch and CW Groetsch. *Inverse problems in the mathematical sciences*, volume 52. Springer, 1993.
- [20] Mario Bertero, Patrizia Boccacci, and Christine De Mol. *Introduction to inverse problems in imaging*. CRC press, 2021.
- [21] Martin Buehrer, Klaas P Pruessmann, Peter Boesiger, and Sebastian Kozerke. Array compression for mri with large coil arrays. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 57(6):1131–1139, 2007.
- [22] Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.
- [23] Julio Martin Duarte-Carvajalino and Guillermo Sapiro. Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization. *IEEE Transactions on Image Processing*, 18(7):1395–1408, 2009.
- [24] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.
- [25] David S Taubman and Michael W Marcellin. Jpeg2000: Standard for interactive imaging. *Proceedings of the IEEE*, 90(8):1336–1357, 2002.
- [26] Lei Yu, Hong Sun, Jean-Pierre Barbot, and Gang Zheng. Bayesian compressive sensing for cluster structured sparse signals. *Signal processing*, 92(1):259–269, 2012.
- [27] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical Image Analysis*, page 102479, 2022.
- [28] Mehmet Akçakaya, Steen Moeller, Sebastian Weingärtner, and Kâmil Uğurbil. Scan-specific robust artificial-neural-networks for k-space interpolation (raki) reconstruction: database-free deep learning for fast imaging. *Magnetic resonance in medicine*, 81(1):439–453, 2019.
- [29] Tianming Du, Honggang Zhang, Yuemeng Li, Stephen Pickup, Mark Rosen, Rong Zhou, Hee Kwon Song, and Yong Fan. Adaptive convolutional neural networks for accelerating magnetic resonance imaging via k-space data interpolation. *Medical image analysis*, 72:102098, 2021.
- [30] Mark Murphy, Marcus Alley, James Demmel, Kurt Keutzer, Shreyas Vasanawala, and Michael Lustig. Fast l_1 -spirit compressed sensing parallel imaging mri: scalable parallel implementation and clinically feasible runtime. *IEEE transactions on medical imaging*, 31(6):1250–1262, 2012.
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.

- [32] Mohammad Zalbagi Darestani and Reinhard Heckel. Accelerated mri with un-trained neural networks. *IEEE Transactions on Computational Imaging*, 7:724–733, 2021.
- [33] Salman UH Dar, Mahmut Yurt, Mohammad Shahdloo, Muhammed Emrullah Ildız, Berk Tınaz, and Tolga Çukur. Prior-guided image reconstruction for accelerated multi-contrast mri via generative adversarial networks. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1072–1087, 2020.
- [34] Alper Güngör, Salman UH Dar, Şaban Öztürk, Yılmaz Korkmaz, Hasan A Bedel, Gokberk Elmas, Muzaffer Ozbey, and Tolga Çukur. Adaptive diffusion priors for accelerated mri reconstruction. *Medical image analysis*, 88:102872, 2023.
- [35] Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, pages 1–9, 2024.
- [36] Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces, 2021.
- [37] Meer Mehran Rashid, Tanu Pittie, Souvik Chakraborty, and NM Anoop Krishnan. Learning the stress-strain fields in digital composites using fourier neural operator. *Iscience*, 25(11), 2022.
- [38] Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- [39] Steven Guan, Ko-Tsung Hsu, and Parag V Chitnis. Fourier neural operator network for fast photoacoustic wave simulations. *Algorithms*, 16(2):124, 2023.
- [40] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations, 2021.
- [41] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations, 2020.
- [42] Miguel Liu-Schiaffini, Julius Berner, Boris Bonev, Thorsten Kurth, Kamyar Azizzadenesheli, and Anima Anandkumar. Neural operators with localized integral and differential kernels. In *Forty-first International Conference on Machine Learning*, 2024.
- [43] Christoph Juchem, Omar M Nahhass, Terence W Nixon, and Robin A de Graaf. Multi-slice mri with the dynamic multi-coil technique. *NMR in Biomedicine*, 28(11):1526–1534, 2015.
- [44] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdal, Adriana Romero, Michael Rabbat, Pascal Vincent, Nafissa Yakubova, James Pinkerton, Duo Wang, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: An open dataset and benchmarks for accelerated MRI, 2018.
- [45] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P Recht, Daniel K Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine*, 79(6):3055–3071, 2018.
- [46] Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on Machine Learning*, pages 399–406, 2010.
- [47] Jian Sun, Huibin Li, Zongben Xu, et al. Deep admm-net for compressive sensing mri. *Advances in neural information processing systems*, 29, 2016.

- [48] Morteza Mardani, Qingyun Sun, David Donoho, Vardan Pappyan, Hatef Monajemi, Shreyas Vasanaawala, and John Pauly. Neural proximal gradient descent for compressive imaging. *Advances in Neural Information Processing Systems*, 31, 2018.
- [49] Jeremy Ocampo, Matthew A Price, and Jason D McEwen. Scalable and equivariant spherical cnns by discrete-continuous (disco) convolutions. *arXiv preprint arXiv:2209.13603*, 2022.
- [50] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023.
- [51] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.
- [52] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *Asilomar Conference on Signals, Systems & Computers*, 2003.
- [53] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.

Appendix

In the appendix, we first present more details of the proposed U-shaped DISCO Neural Operator (UDNO, in Section A). We then provide more details of the machine learning framework implementation (Section B) as well as additional numerical and visualization results of the multi-pattern and multi-rate subsampling experiments (Section C). Finally in Section D we show that our model performance improves when increasing the number of cascades of the unrolled layers (the INO and data consistency layer in Fig. 2).

A UDNO Architecture

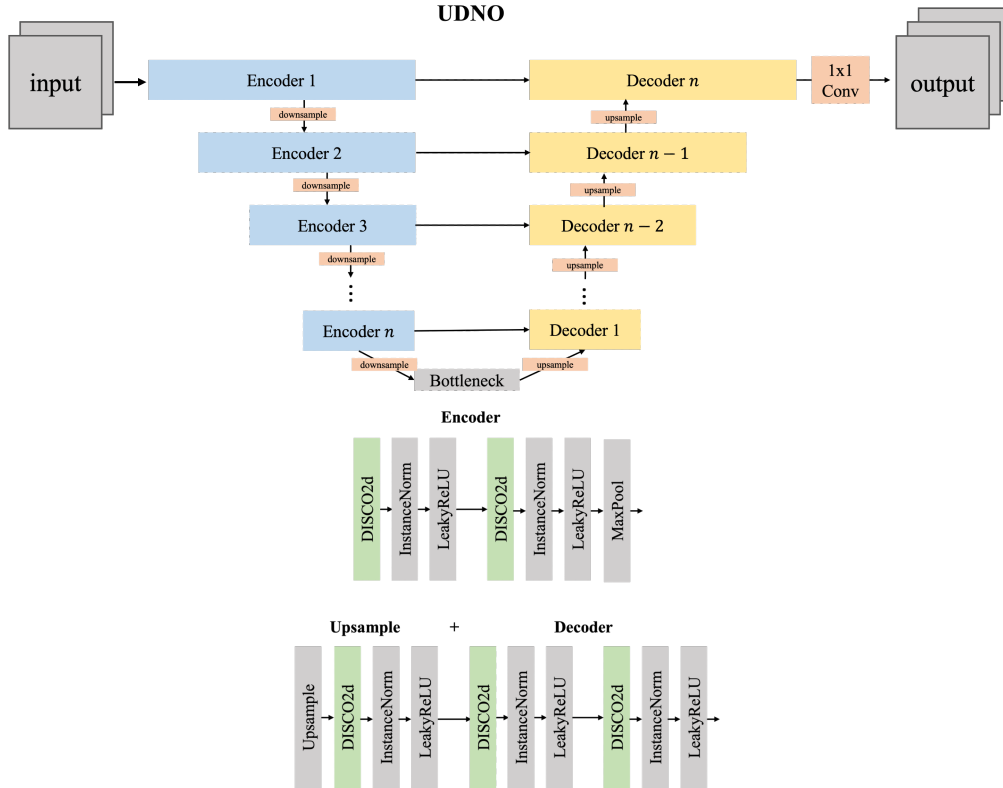


Figure 4: **UDNO architecture.** We propose a U-shaped neural operator UDNO to capture multi-scale features of the input. The UDNO uses discrete-continuous convolutions (DISCOs) [49] as the local integral operator. The final 1×1 convolution allows the module to flexibly project to the desired number of output channels and is resolution invariant by virtue of being a pointwise operation. The UDNO is an end-to-end *neural operator*.

We provide additional details of the proposed UDNO (U-Shaped DISCO Neural Operator) architecture. Fig. 4 depicts the overall architecture, which mimics the U-Net [53]. We use the updated implementation of the U-Net in [1]. Our network architecture has two differences. First, all traditional convolutions are replaced with their DISCO counterparts. Second, transpose convolutions are replaced by an interpolation upsampling step, followed by a DISCO2d layer, InstanceNorm layer, and LeakyReLU activation. DISCO2d layers function as drop-in replacements for traditional 2d convolution layers. They do not change the spatial dimension of the input. The UDNO is an end-to-end *neural operator* model.

As in the traditional U-Net [53], each encoder block halves the spatial dimensions and doubles the feature channels. Each decoder step (upsampling + decoder) doubles the spatial dimensions and halves the feature channels. Skip connections are included, as in the original architecture. All




Figure 5: Subsampling mask patterns. The visualized patterns are all for the 4x acceleration rate. **Top:** Rectilinear patterns. Equispaced, Random, Magic. **Bottom:** Irregular patterns. Gaussian, Radial, Poisson

Table 1: k -space subsampling configurations (acceleration and center fraction parameters) used for MRI experiments.

Alias	Acceleration rate	Center fraction rate
16x	16	0.02
8x	8	0.04
6x	6	0.06
4x	4	0.08

components of the UDNO operate in the function space and are not tied to a specific discretization, thus making the model an *end-to-end* neural operator.

B Additional Implementation Details

We summarize the configurations of different CS-MRI subsampling rates in Table 1 and subsampling patterns in Fig. 5.

B.1 Learning Sensitivity Maps for Multi-Coil MRI

In reconstruction, the sensitivity map S_i for the i^{th} coil is needed for coil reductions and expansions. Inspired by [1], we use a UDNO with 4 encoder/decoder steps, 8 hidden channels, 0.02 DISCO radius (assuming the domain is $[-1, 1]^2$), and the kernel basis from [42] with 1 isotropic basis and 5 anisotropic basis rings, each containing 7 basis functions. We use this UDNO to predict the sensitivity map S_i from the input coil measurement \mathbf{k}_i . We then follow [1] to combine multiple coils weighted by the corresponding learned sensitivity maps.

B.2 UDNO and DISCO Details

Both the KNO and INO use DISCO layers using the kernel basis from [42] with 1 isotropic basis and 5 anisotropic basis rings, each containing 7 basis functions. The KNO (measurement space neural operator) is implemented as a UDNO with 2 input and output channels, 16 hidden channels, and 4 depth (encoder/decoder steps). KNO DISCO kernels have a radius cutoff of 0.04. The INO (image-space neural operator) is implemented as a UDNO with 2 input and output channels, 18 hidden channels, and 4 encoder/decoder steps. INO DISCO kernels have a smaller radius cutoff of 0.02 with the same internal basis shape. We train both our model and the baseline with SSIM loss, 0.0003 learning rate, and 16 batch size on 4 A100 (40G) GPUs.

C Additional Multi-Pattern and Multi-Rate Results

We summarize the numerical results of the proposed neural operator and the end-to-end VarNet baseline [1]: Table 2 for different subsampling patterns and Table 3 for different subsampling rates. We also visualize the numerical performance in Fig.6, where we observe that, unlike the baseline method, the proposed neural operator achieves consistent performance across different subsampling patterns and rates.

We also present the visual comparisons in Fig.7, where we see the neural operator outperforms the baseline.

Table 2: Performance comparison of NO-VN (ours) and E2E-VN [1] (baseline) models across different subsampling patterns. The metrics used are PSNR (dB), SSIM, and NMSE.

Pattern	NO-VN (Ours)			E2E-VN (Baseline)		
	PSNR \uparrow	SSIM \uparrow	NMSE \downarrow	PSNR \uparrow	SSIM \uparrow	NMSE \downarrow
Equispaced	37.219	0.925	0.011	37.264	0.917	0.008
Random	36.155	0.914	0.015	37.059	0.913	0.008
Magic	37.635	0.935	0.011	37.639	0.927	0.008
Gaussian	37.286	0.908	0.009	32.246	0.868	0.037
Radial	37.640	0.922	0.008	34.163	0.897	0.015
Poisson	39.683	0.944	0.005	35.393	0.922	0.011

Table 3: Performance comparison of NO-VN (ours) and E2E-VN (baseline) models across different subsampling rates. The metrics used are PSNR (dB), SSIM, and NMSE.

Rate	NO-VN (Ours)			E2E-VN (Baseline)		
	PSNR \uparrow	SSIM \uparrow	NMSE \downarrow	PSNR \uparrow	SSIM \uparrow	NMSE \downarrow
2x	39.762	0.942	0.005	39.825	0.938	0.005
4x	36.959	0.899	0.008	36.400	0.893	0.009
6x	35.149	0.873	0.012	34.630	0.867	0.014
8x	33.712	0.854	0.017	33.087	0.846	0.019

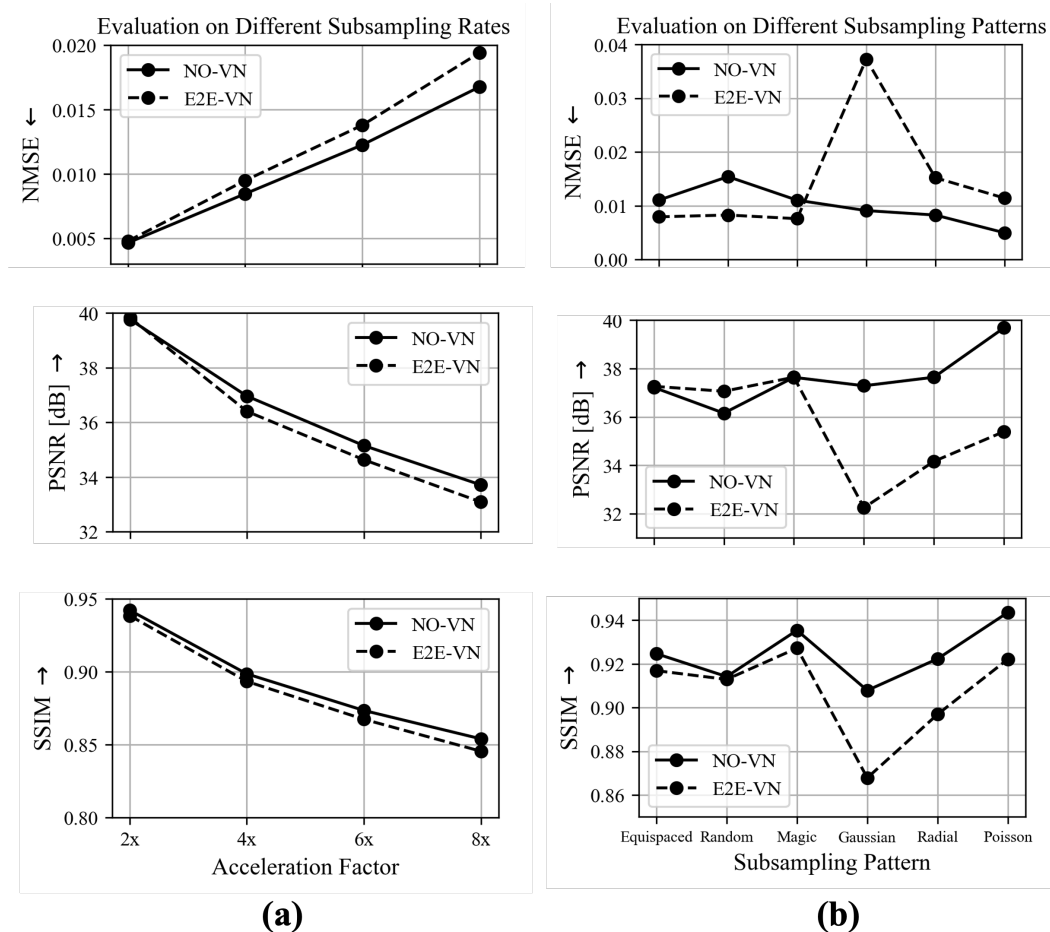


Figure 6: NMSE(\downarrow), PSNR(\uparrow), and SSIM(\uparrow) results for NO-VN (ours) and E2E-VN (SOTA). **Left:** Multi-rate. **Right:** Multi-pattern.

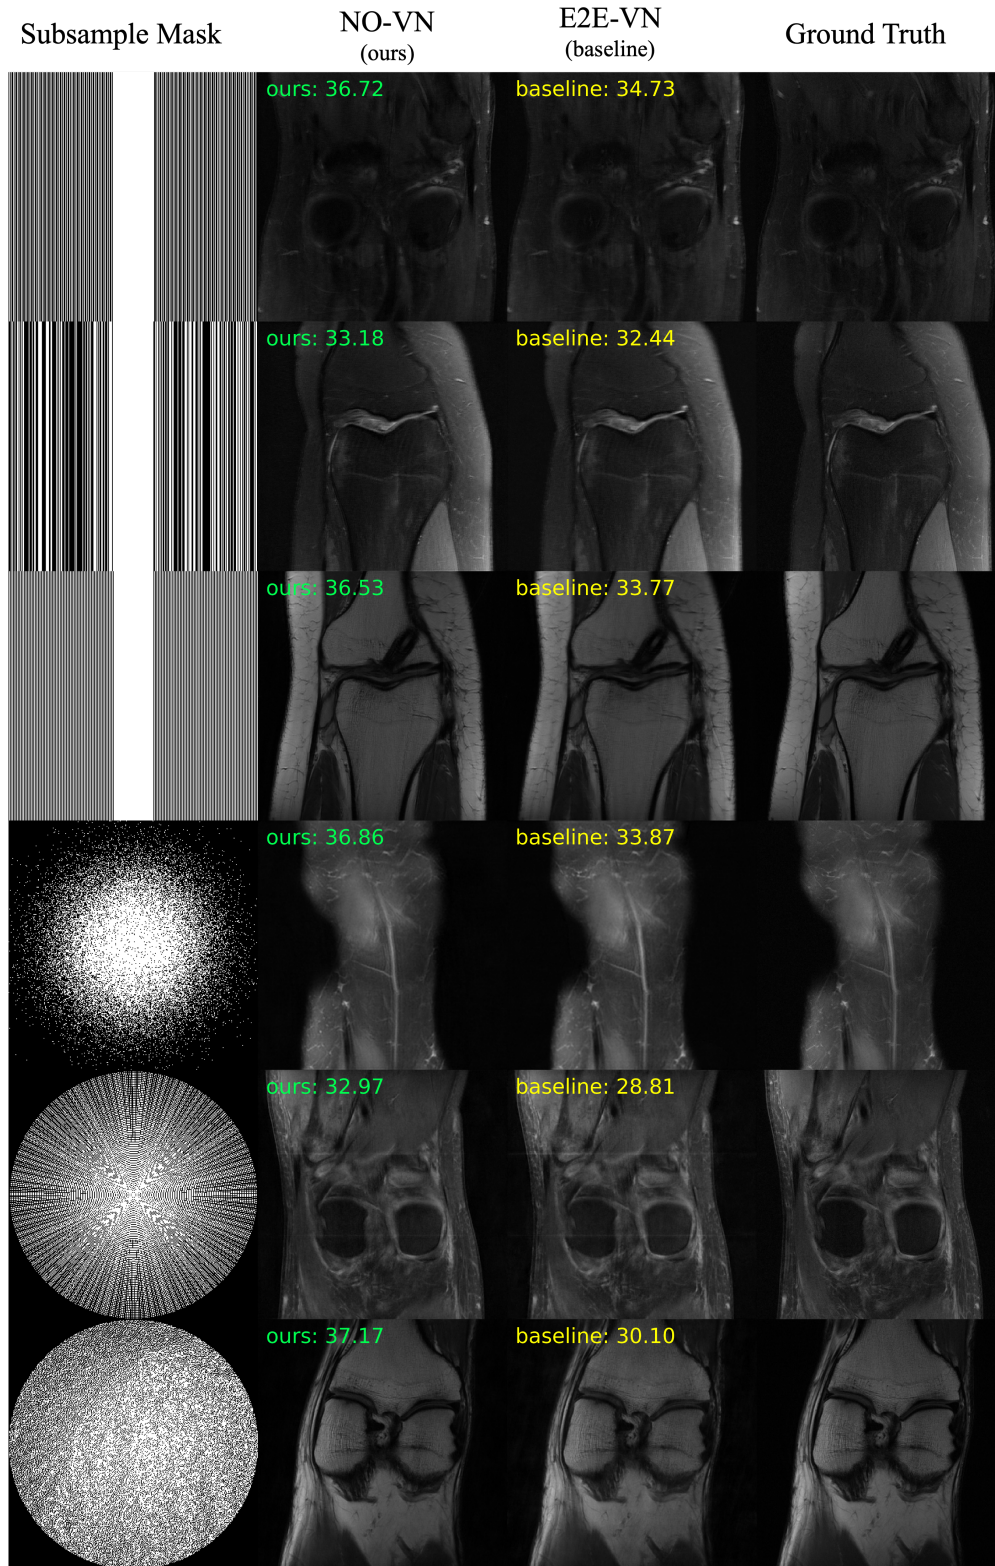


Figure 7: **Visualizations.** Inference comparisons between our multi-pattern models, each trained under identical conditions. **Left:** NO-VN (our model). **Right:** current SOTA baseline E2E-VN. The rows represent different subsampling mask patterns, listed from top to bottom: equispaced, random, magic, Gaussian, radial, and Poisson. All subsampling is performed at a 2x rate, and the reconstruction quality is measured using PSNR.

D Scaling with Cascades

We train and evaluate our NO-VN neural operator model on single resolution 4x equispaced subsampled data, at increasing number of cascades. We find that as we increase the number of cascades, reconstruction performance improves, as is expected in unrolled neural networks (Table 4 and Fig.8).




Figure 8: PSNR \uparrow , SSIM \uparrow , and NMSE \downarrow plotted against number of cascades. As we increase the number of cascades, reconstruction performance improves across all metrics.

Table 4: Performance comparison across different numbers of cascades. The metrics used are PSNR (dB), SSIM, and NMSE.

Cascades	PSNR \uparrow	SSIM \uparrow	NMSE \downarrow
1	37.008	0.898	0.0084
4	37.657	0.905	0.0076
8	38.158	0.909	0.0069